A MARKET DRIVER VOLATILITY MODEL VIA POLICY IMPROVEMENT ALGORITHM

JUN MAEDA AND SAUL D. JACKA

Abstract. In the over-the-counter market in derivatives, we sometimes see large numbers of traders taking the same position and risk. When there is this kind of concentration in the market, the position impacts the pricings of all other derivatives and changes the behaviour of the underlying volatility in a nonlinear way.

We model this effect using Heston's stochastic volatility model modified to take into account the impact. The impact can be incorporated into the model using a special product called a market driver, potentially with a large face value, affecting the underlying volatility itself. We derive a revised version of Heston's partial differential equation which is to be satisfied by arbitrary derivatives products in the market. This enables us to obtain valuations that reflect the actual market and helps traders identify the risks and hold appropriate assets to correctly hedge against the impact of the market driver.

arXiv:1612.00780v1 [q-fin.PR] 2 Dec 2016

1. Introduction

Japan has one of the largest equity derivatives markets in the world. According to the Bank for International Settlements, Japan had $378 billion in face value of equity-linked contracts out of the worldwide total of $5,445 billion as of September 13, 2015 [1]. A common underlying for equity-linked derivatives products in the country is the price-weighted Nikkei Stock Average Index (Nikkei 225) published by Nikkei Inc. Since the country is the world's third largest economy by GDP, people generally assume that the market is liquid enough to trade freely any desired position. However, in the first author's experience, this is not quite true. Long-dated volatility (especially between 2 and 5 years) is generally priced low due to the fact that the most of the traders in the market already own vega (sensitivity to volatility) from selling (usually in significant sizes) a structured product called an 'auto-callable' to their clients. Therefore it is generally difficult to sell vega in the market. We will discuss this exotic case in more detail in a future paper, and for now, we will focus on explaining our model in a simpler context in this paper. The important fact to note is that there is a position that affects the pricings and risks of all the existing and potential derivatives products in the market.

In order to understand the background of our model, we first formulate a toy example.

Assume that there are only 2 traders, A and B, in the over-the-counter (OTC) market. If A wants to buy volatility, A needs to buy it from B and vice versa. If A buys $10 million of vega

Date: December 5, 2016.

Key words and phrases. stochastic volatility model, Heston model, semilinear parabolic partial differential equations, policy improvement algorithm, Hamilton-Jacobi-Bellman equation.

from B, the vega that B holds is decreased by the same amount. Generally, traders don't want to own so much risk on one side, so they might want to hedge the risk a little. Since the only market participants are A and B, they need to reverse what they previously traded in order to hedge themselves. This does not make much sense in this case, as there are only 2 market participants, but even if we assumed more participants in the market, this is still essentially what is happening: overall, the market vega is maintained and does not change whatever A and B do. Whatever A gains, B loses and vice versa.

Now introduce a new market participant C. Let us assume that C is not a participant in the OTC market but only buys vega from A and B as their client to hedge against market risk, and doesn't otherwise hedge the position (we could think of C as an insurance company, for example). If C buys \$10 million of vega from A, then A is now short the risk, so may want to buy some back in the market to hedge himself. A needs to buy it from B, of course, as C doesn't sell any vega. The important point is that the OTC market whose only participants are A and B is now short \$10 million of vega overall. The market now would like to buy some vega back. This generally drives the volatility of the underlying security or index higher.

We elaborate this point in more detail. The demand and supply of vega could be, in general, directly converted to the supply and demand of volatility. It is easier to think of this in the Black-Scholes framework. If there is more demand for vega than supply, more people want to buy vega. The way they accomplish this is to buy plain vanilla calls and puts, which are positive vega products. If more people buy these products, the prices of the products move higher. Given other parameters are fixed, this price increase could only be explained by the increase in the underlying volatility. This is why the actual market participants refer to 'buying (selling) volatility' when they are actually buying (selling) vega. These phrases will be used with the same meanings hereafter.

Remark 1. The corresponding volatility level is implied volatility as opposed to realized (or historical) volatility.

Up to this point, volatility movement is just a matter of demand and supply. Now suppose that the derivative product that C bought has big second order risks, like vanna (the derivative of vega with respect to stock price) and volga (the derivative of vega with respect to volatility). For example, if the product is long vanna, vega increases when the underlying stock moves higher. In this case, A gets shorter vega just from the market movement and he needs to buy it in the market to rehedge himself. However, if B has the same position, B gets shorter vega as well, so neither of them are interested in selling any more vega. This will make the volatility even higher. Note that in this situation, what is moving the volatility is just the change in the risk of the product that was already traded, not a new trade. We call the special product (of which the risks affect the dynamics of supply and demand of the volatility) the *market driver*.

In order to model the example above, we posit a simple and easy-to-use model which is just an extension of the Heston model, one of the most popular stochastic volatility models. The core of our model is a semilinear parabolic partial differential equation (PDE) that we retrieve to price the

market driver. Once we obtain the valuation of the market driver, we use a *linear* parabolic PDE, which is very similar to those of Black-Scholes and Heston, to price other derivatives products.

As mentioned earlier, we are more interested in the case where the market driver is of a specific exotic type because we think its risk feedback effect is more prominent in practice. We will handle this problem in a future paper and concentrate now on the case when it is of plain vanilla type.

The problem statement so far may remind some readers of the 'feedback effect' of options which is now a somewhat mature field. The Black-Scholes model with a feedback effect models the prices of derivative products affected by delta hedging executed by program traders [5, 19, 21]. It was a field which attracted a lot of interests in the 1990s. However, as far as we know, not much work has been done since then. Although the research in this paper was done separately from the studies done in the field, our ideas are very similar in the sense that some trade affects other option pricing. We are (in a way) incorporating the feedback effect in a stochastic volatility framework. The key difference, however, is that we are not applying the feedback effect of the underlying asset (stock), but instead, *that of the underlying volatility*. In the earlier models, the effect impacts the volatility passively via program traders trading the underlying asset. On the other hand, our model incorporates the effect directly in the dynamics of the volatility.

It may not look natural to incorporate a feedback effect in the volatility as it is not a tradable asset. However, from the first author's experience, supply and demand effects of the volatility do exist in the actual market and we think our model reflects, at least qualitatively, the actual market dynamics with the market driver. We believe that our model is more in line with market practitioners' perspectives than the classical feedback model.

One of the difficulties in the earlier feedback models is that they model the realized (historical) volatility rather than the implied volatility. Hedging delta of derivative products by dynamically trading the underlying asset does affect the implied volatility, but only that of short maturity. Behaviour of the current stock price has little impact over the long-dated implied volatility. Depending on the sign of the vanna of the market driver, it is possible, for example, that even when the realized volatility increases with a large drop in the stock price, the long-dated implied volatility go lower. This cannot be modelled in the classic feedback model.

One of the benefits of our model is that the nonlinear PDEs that we derive can be approximated by a series of linear ones. The PDEs derived in the classic feedback model are generally of quasilinear type, where the nonlinearity occurs in the highest order of the equations. On the other hand, although we need a pair of PDEs, one for the market driver and the other for a general derivative product, our PDEs are at most of semilinear parabolic type, where the nonlinearity occurs in lower order of the equations. This enables us to apply a linear approximation algorithm called the Policy Improvement Algorithm (PIA) in which the approximated solution converges quickly to the actual solution of the semilinear PDE.

The reason why we introduce the PIA is that it enables us to reuse the setup for the Heston model. The Heston model has already been implemented in practice and is widely used. It is

convenient to use the existing setup, whenever possible, to calculate the solutions of the new model. We also note that in the course of our research, we encountered some cases where we had a convergence in numerical solution using the PIA, but not using the finite difference method (FDM): the PIA seems to have better convergence in numerical solution than the FDM.

The rest of the paper is organized as follows: Section 2 explains the new model in detail. We will establish the existence and uniqueness of the solution to our PDEs in Section 3. In Section 4, we transform the nonlinear PDE to an HJB equation. The PIA is then described in Section 5. In Section 6, we give a numerical example to see how valuations and risks, which are very important for day-to-day hedging for traders, change in our model from those in Heston's model. We also see in this section how PIA-approximated solutions converge to that of the nonlinear PDE. We give our conclusions in Section 7.

2. The Market Driver Model

We start by briefly reviewing Heston's stochastic volatility model [7]. The stochastic differential equations for the stock price and the variance are:

$$\begin{cases} dS = \mu S dt + \sqrt{v} S dW^1 \\ dv = \kappa(\bar{v} - v)dt + \eta\sqrt{v} dW^2 \\ \langle dW^1, dW^2 \rangle = \rho dt. \end{cases} \tag{2.1}$$

Here, S denotes the underlying stock price and v the variance of the underlying. W^1 and W^2 are Wiener processes with correlation ρ, μ is the drift of the stock, $\kappa > 0$ is a constant which expresses the intensity of the mean reversion of the variance, $\bar{v}$ is the mean variance, and η is the volatility of the variance.

Since v only takes positive values, we usually require the model to satisfy Feller's condition for avoiding the origin [13]:

$$2\kappa\bar{v} > \eta^2. \tag{2.2}$$

With this setup, the value V of a derivatives product with its satisfies Heston's PDE:

$$\frac{\partial V}{\partial t} + rS\frac{\partial V}{\partial S} + \kappa\Big(\bar{v} - \omega v\Big)\frac{\partial V}{\partial v} + \frac{1}{2}vS^2\frac{\partial^2 V}{\partial S^2} + \frac{1}{2}v\eta^2\frac{\partial^2 V}{\partial v^2} + vS\eta\rho\frac{\partial^2 V}{\partial S \partial v} - rV = 0 \tag{2.3}$$

with appropriate initial (or terminal, if we are calculating backwards in time) and boundary conditions. Equation (2.3) is a second order linear parabolic PDE. Here, ω is some constant for volatility risk premium and r is the interest rate.

Let us now assume that there is some distinguished product (called the market driver) with value denoted by F.

Using this F, our revised model is written as

$$\begin{cases} dS = \mu S dt + \sqrt{v} S dW^1 \\ dv = \kappa(\bar{v} - v + Q\frac{\partial F}{\partial v})dt + \eta\sqrt{v}dW^2 \\ d\langle W^1, W^2\rangle_t = \rho dt \end{cases} \tag{2.4}$$

with some coefficient Q.

Note that the only change made to the Heston SDE (2.1) is the term $\kappa Q\frac{\partial F}{\partial v}$ in the second equation. A simple justification for this is that the vega (in this paper, we use the term 'vega' for the derivative of the valuation with respect to variance, whereas it usually means the derivative of the valuation with respect to volatility) of the market driver impacts supply and demand of the variance and causes the shift in its mean. We are only adding this adjustment to the variance SDE. If we want to, we could, of course, similarly add 'delta' (derivatives of valuation with respect to the underlying stock price) adjustment in the SDE for the stock price S in (2.4). However, we do not do this since i) deltas of derivatives products are generally low, so in order to have a large impact on the stock price, the face value traded on the position needs to be massive, which is not realistic and ii) the stock market is more liquid than the OTC derivatives market, in the sense that there are more people with different incentives in trading and many more people have access to the market (for example, personal investors can easily trade stocks, whereas they might need to satisfy additional requirements in order to trade derivatives. It is even harder for them to be able to trade in the OTC market due to size requirements, credit issues, and other restrictions).

A sufficient condition for the variance not to go negative is derived by comparing the 2 processes v and v' starting at the same value:

$$\begin{cases} dv = \kappa(\bar{v} - v + Q\frac{\partial F}{\partial v})dt + \eta\sqrt{v}dW^2 \\ dv' = \kappa\{\bar{v} - v' + \min(Q\frac{\partial F}{\partial v})\}dt + \eta\sqrt{v'}dW^2. \end{cases} \tag{2.5}$$

Since we will be working in a bounded domain, Proposition 5.2.18 in [12] shows that $v' \leq v$ almost surely. Applying Feller's condition (2.2) on v', if

$$2\kappa\Big\{\bar{v} + \min\Big(Q\frac{\partial F}{\partial v}\Big)\Big\} > \eta^2, \tag{2.6}$$

then $v' > 0$ almost surely, hence $v > 0$ almost surely. We call condition (2.6) the *positive variance condition*.

If we follow the usual argument, we obtain the following PDE for the value V of a derivative:

$$\frac{\partial V}{\partial t} + rS\frac{\partial V}{\partial S} + \kappa\Big(\bar{v} - \omega v + Q\frac{\partial F}{\partial v}\Big)\frac{\partial V}{\partial v} + \frac{1}{2}vS^2\frac{\partial^2 V}{\partial S^2} + \frac{1}{2}v\eta^2\frac{\partial^2 V}{\partial v^2} + vS\eta\rho\frac{\partial^2 V}{\partial S\partial v} - rV = 0. \tag{2.7}$$

Since F is also the value of a specific derivative, we can substitute $V = F$ in (2.7) and obtain a nonlinear PDE for F:

$$\frac{\partial F}{\partial t} + rS\frac{\partial F}{\partial S} + \kappa\Big(\bar{v} - \omega v + Q\frac{\partial F}{\partial v}\Big)\frac{\partial F}{\partial v} + \frac{1}{2}vS^2\frac{\partial^2 F}{\partial S^2} + \frac{1}{2}v\eta^2\frac{\partial^2 F}{\partial v^2} + vS\eta\rho\frac{\partial^2 F}{\partial S\partial v} - rF = 0. \tag{2.8}$$

Note that given F, the differential equation (2.7) is a second order parabolic PDE that is linear in V as in the Heston model. On the other hand, the differential equation (2.8), is semilinear.

3. Partial Differential Equations

We recall some theorems from the theory of PDEs. For more detail, refer to [14].

We take a bounded, open, and connected domain $\mathcal{E}$ in $\mathbb{R}^2_+$ which is bounded away from the axes. We further assume that $\partial\mathcal{E}$ is $C^{2+\alpha'}$ for some $\alpha' > 0$. Let $Q_T = \mathcal{E} \times (0,T)$, $\mathcal{D} = \partial\mathcal{E}$, $\mathcal{D}_T = \{(x,y,t)|(x,y) \in \mathcal{D}, t \in [0,T]\}$, and $\Gamma_T = \mathcal{D}_T \cup \{(x,y,t)|(x,y) \in \mathcal{E}, t = 0\}$. We impose ψ as our initial and boundary conditions and assume it satisfies the compatibility condition, i.e. $\psi(x,y,t) \in C(\overline{Q_T})$.

We define the differential operator L by

$$\begin{aligned} -Lu :&= rxu_x + \kappa(v_0 - \alpha y)u_y + \frac{1}{2}x^2 y u_{xx} + \frac{1}{2}\eta^2 y u_{yy} + \eta\rho x y u_{xy} \\ &= a_{ij}u_{ij} + b_i u_i \end{aligned} \tag{3.1}$$

under the Einstein summation convention.

We reparameterize time-to-go t backwards by replacing $t \to T - t$ and rewrite (2.8) in general form:

$$u_t + Lu + ru - \kappa Q {u_y}^2 = 0. \tag{3.2}$$

The PDE (3.2) is uniformly parabolic as it satisfies

$$\nu_1|\xi|^2 \le a_{ij}\xi_i\xi_j \le \nu_2|\xi|^2 \quad \forall (x,y) \in \overline{\mathcal{E}}, \quad \forall \xi \in \mathbb{R}^2 \tag{3.3}$$

for some $\nu_1, \nu_2 > 0$.

Theorem 6.2 of Chapter V of [14] shows the existence and uniqueness of the solution to (3.2) with continuous initial and boundary conditions. By the theorem, the solution belongs to the space $H^{\beta,\beta/2}(\overline{Q_T})$ for some $0 < \beta < 1$, it also has bounded first spatial derivatives in $\overline{Q_T}$, and its second order spatial derivatives and first order time derivative belong to $H^{\gamma,\gamma/2}(\overline{Q_T})$ for some nonnegative and nonintegral number γ.

By substituting this solution in the coefficients of the PDE (2.7), Corollary 1 in Section 3.5 on page 74 of [6] affirms the existence and uniqueness of the solution to the linear PDE for suitable initial and boundary conditions.

Remark 2. We require the positive variance condition (2.6) to be satisfied in order to ensure that v is nonnegative. Theorem 6.2 of Chapter V from [14] affirms that F_y is bounded, but as far as the statement of the theorem goes, we don't have an explicit expression of it. For that reason, it is not easy to show that (2.6) is satisfied in general. In the case where the market driver with value F is of plain vanilla type with $Q > 0$, enforcing Feller's condition (2.2) is sufficient for the positive variance condition (2.6) to be satisfied since $\partial F/\partial y \geq 0$, and therefore $Q(\partial F/\partial y) \geq 0$.

4. Control Problem

From now on, we focus on solving (3.2). We can apply various numerical methods, for example, the FDM, to calculate the solution numerically. If we were to do this, we would need additional resources to implement it in actual trading and in some cases, it may not be easy to do so. One of the difficulties may originate from the fact that even though it's semilinear, it's still a nonlinear PDE that we are dealing with. Applying the model to actual trading becomes more straightforward with a help of the Policy Improvement Algorithm (PIA).

It is easy to see that (3.2) can be rewritten as

$$\inf_{\pi\in\mathbb{R}}\left(u_t + Lu + ru - \pi u_y + \frac{\pi^2}{4\kappa Q}\right) = 0. \tag{4.1}$$

Note that this is the HJB equation to minimize

$$V^\pi(x,y,t) = \mathbb{E}\left[\int_0^{\tau\wedge t} e^{-rs} f^\pi(Z_s^{z,\pi}, t-s)ds + e^{-r(\tau\wedge t)} g(Z_{\tau\wedge t}^{z,\pi}, t\wedge\tau)\right] \tag{4.2}$$

under the controlled process $Z_t^{z,\pi} := (X, Y^\pi)^T$ with SDEs

$$\begin{cases} dX = \mu X dt + \sqrt{Y} X dW^1 \\ dY^\pi = \kappa(\bar{v} - Y^\pi + \pi/\kappa)dt + \eta\rho\sqrt{Y^\pi}dW^1 + \eta\sqrt{Y^\pi}\sqrt{1-\rho^2}dW^2 \\ d\langle W^1, W^2\rangle_t = 0 \end{cases} \tag{4.3}$$

with $Z_0^{z,\pi} = z = (x,y)^T$. Here, $f^\pi = \pi^2/4\kappa Q$, $g = \psi$ is the initial and boundary conditions introduced in Section 3, and τ is the first hitting time of the boundary of the domain.

Our problem is now converted into the HJB equation for the following controlled initial/boundary problem:

$$\begin{cases} \inf_{\pi\in\mathbb{R}} \left(u_t + Lu + ru - \pi u_y + \dfrac{\pi^2}{4\kappa Q} \right) = 0 \quad (x,y,t) \in \mathcal{E} \times (0,T) \\ u(x,y,t) = \inf_{\pi} V^{\pi}(x,y,t). \end{cases} \tag{4.4}$$

From the positive variance condition (2.6),

$$\pi > \frac{\eta^2}{2} - \kappa\bar{v} \tag{4.5}$$

is sufficient for Y not to go below zero.

5. Policy Improvement Algorithm

We now give a detailed formulation of the PIA and the proof of convergence. For more detail, refer to [10] and [11].

Let $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$ be a filtered probability space satisfying the usual conditions that supports a 2-dimensional $(\mathcal{F}_t)_{t\geq 0}$ - Wiener process $W = (W_t)_{t\geq 0}$.

For any process $\mathcal{Y} = (\mathcal{Y}_t)_{t\geq 0}$, define

$$\tau_{\mathcal{E}}(\mathcal{Y}) := \inf\{t \geq 0; \mathcal{Y}_t \in \partial\mathcal{E}\}. \tag{5.1}$$

Let

$$\begin{aligned} \mathcal{A}(z,T) := \{\Pi = (\Pi_t)_{t<T}; \Pi \text{ is adapted to } (\mathcal{F}_t)_{t<T}, \Pi_t(\omega) \in \mathbb{R} \text{ for every } t < T \text{ and } \omega \in \Omega, \\ \text{and there exists a process } Z^{z,\Pi} \text{ that satisfies (5.3) and is unique in law}\}, \end{aligned} \tag{5.2}$$

where

$$Z_t^{z,\Pi} = z + \int_0^t \sigma(Z_s^{z,\Pi}, s, \Pi_s) dW_s + \int_0^t \mu(Z_s^{z,\Pi}, s, \Pi_s) ds \qquad t \leq T \wedge \tau_{\mathcal{E}}(Z^{z,\Pi}). \tag{5.3}$$

A measurable function $\pi : \Omega \times (0,T) \to \mathbb{R}$ is a *Markov policy* if for every $z \in \mathcal{E}$ and $T > 0$ there exists a process $Z_t^{z,\pi}$ that is unique in law and satisfies the following:

$$\begin{aligned} Z_t^{z,\pi} &= z + \int_0^t \sigma(Z_s^{z,\pi}, s, \pi(Z_s^{z,\pi}, s)) dW_s + \int_0^t \mu(Z_s^{z,\pi}, s, \pi(Z_s^{z,\pi}, s)) ds \\ &= z + \int_0^t \sigma_\pi(Z_s^{z,\Pi}, s) dW_s + \int_0^t \mu_\pi(Z_s^{z,\Pi}, s) ds \qquad t \leq T \wedge \tau_{\mathcal{E}}(Z^{z,\pi}). \end{aligned} \tag{5.4}$$

For any domain $Q_T = \mathcal{E} \times (0,T)$ and bounded measurable function g defined on Γ_T, define $V^{g,\mathcal{E},\pi}$ by

$$V^{g,\mathcal{E},\pi}(z,t) = \mathbb{E}_z\bigg(\int_0^{t\wedge\tau} e^{-rs} f^\pi(Z_s^{z,\pi}, t-s)ds + e^{-r(t\wedge\tau)} g(Z_{t\wedge\tau}^{z,\pi}, t\wedge\tau)\bigg), \tag{5.5}$$

where f^π is the running cost and τ is the first exit time from Q_T.

Now define

$$V^{g,\mathcal{E}} := \inf_{\pi\in\mathcal{A}} V^{g,\mathcal{E},\pi}. \tag{5.6}$$

Finally, we define the differential operator L^π:

$$L^\pi u := -u_t + \frac{1}{2} Tr\{\sigma_\pi^T (Hu)\sigma_\pi\} + \mu_\pi^T \nabla u \quad \text{ for } u \in C^{2,1}, \tag{5.7}$$

where Hu is the Hessian of the function u.

Proposition 5.1. *For any Markov policy π that is Lipschitz on compacts in $\mathbb{R}^2_+$, the following holds: $V^{g,\mathcal{E},\pi} \in C^{2,1}(Q_T)$ and it satisfies*

$$L^\pi V^{g,\mathcal{E},\pi} - rV^{g,\mathcal{E},\pi} + f^\pi = 0. \tag{5.8}$$

Proof. It suffices to prove that $V^{g,\mathcal{E},\pi}$ satisfies (5.8) in every open ball U with $\overline{U} \subset Q_T$. Let U be such an open ball with centre ζ and radius ℓ. Let $z \in U$ and define τ as the first time the process $Z^{z,\pi}$ hits the boundary of U. For every $n \in \mathbb{N}$, define U_n as the closed ball with centre ζ and radius $\ell - \frac{1}{n}$, and let τ_n be the first time the process $Z^{z,\pi}$ hits the boundary of U_n.

Let $v \in C^{2,1}(U) \cap C(\overline{U})$ be the unique solution of the initial boundary value problem

$$\begin{cases} L^\pi v - rv + f^\pi = 0 \\ v|_{\mathcal{D}_T} = V^{g,\mathcal{E},\pi}|_{\mathcal{D}_T}. \end{cases} \tag{5.9}$$

The existence and uniqueness is guaranteed by Corollary 1 on page 71 in [6] and Lemma A.5 in the Appendix. The partial derivatives of v are Hölder continuous by the same corollary. Let n_0 be large enough such that $z \in U_n$, and for every $n \geq n_0$, define the process $(J^n)_{n\geq n_0}$ by

$$J_t^n := \int_0^{t\wedge\tau_n} e^{-rs} f^\pi(Z_s^{z,\pi}, t-s)ds + e^{-r(t\wedge\tau_n)} v(Z_{t\wedge\tau_n}^{z,\pi}, t - t\wedge\tau_n) \tag{5.10}$$

and

$$J_t := \int_0^{t\wedge\tau} e^{-rs} f^\pi(Z_s^{z,\pi}, t-s)ds + e^{-r(t\wedge\tau)} v(Z_{t\wedge\tau}^{z,\pi}, t - t\wedge\tau). \tag{5.11}$$

Ito's formula on $[0, \tau_n]$ and the differential equation for v yield

$$
\begin{aligned}
J^n_t &= v(z,t) + \int_0^{t\wedge\tau_n} e^{-rs}(f^\pi - rv + L^\pi v)(Z^{z,\pi}_s, t-s)ds + \int_0^{t\wedge\tau_n} e^{-rs}(\nabla v)^T \sigma_\pi dW_s \\
&= v(z,t) + \int_0^{t\wedge\tau_n} e^{-rs}(\nabla v)^T \sigma_\pi dW_s.
\end{aligned} \tag{5.12}
$$

Hence J^n is a local martingale, and since it is clearly a bounded process, it is a uniformly integrable martingale. Thus the Dominated Convergence Theorem yields

$$
v(z,t) = \lim_{n\to\infty} E(J^n_0) = \lim_{n\to\infty} E(J^n_t) = E(J_t). \tag{5.13}
$$

From the initial and boundary conditions for v, we obtain

$$
\begin{aligned}
J_t &= \int_0^{t\wedge\tau} e^{-rs} f^\pi(Z^{z,\pi}_s, t-s)ds + e^{-r(t\wedge\tau)} v(Z^{z,\pi}_{t\wedge\tau}, t - t\wedge\tau) \\
&= \int_0^{t\wedge\tau} e^{-rs} f^\pi(Z^{z,\pi}_s, t-s)ds + e^{-r(t\wedge\tau)} V^{g,\mathcal{E},\pi}(Z^{z,\pi}_{t\wedge\tau}, t - t\wedge\tau) \\
&= E\bigg(\int_0^{t\wedge\tau} e^{-rs} f^\pi(Z^{z,\pi}_s, t-s)ds + e^{-r(t\wedge\tau)} g(Z^{z,\pi}_{t\wedge\tau}, t\wedge\tau) \bigg| \mathcal{F}_\mathcal{S} \bigg).
\end{aligned} \tag{5.14}
$$

The last equality in (5.14) follows from Lemma A.1. We conclude:

$$
\begin{aligned}
v(z,t) &= E\bigg(\int_0^{t\wedge\tau} e^{-rs} f^\pi(Z^{z,\pi}_s, t-s)ds + e^{-r(t\wedge\tau)} g(Z^{z,\pi}_{t\wedge\tau}, t\wedge\tau) \bigg) \\
&= V^{g,\mathcal{E},\pi}(z,t)
\end{aligned} \tag{5.15}
$$

♢

We now describe the algorithm. Let π_0 be a Markov policy that is Lipschitz on compacts in $\mathbb{R}^2_+$. The algorithm is defined as follows:

$$
\begin{cases}
\mathcal{L}u_i - \pi_i(u_i)_y + f^{\pi_i} = 0 \\
\pi_{i+1}(z,t) = \arg\min_{p\in A}(L^p u_i(z,t) - r u_i(z,t) + f^p(z,t)),
\end{cases} \tag{5.16}
$$

where the differential operator $\mathcal{L}$ is defined as

$$
\mathcal{L} := -\frac{\partial}{\partial t} - L \tag{5.17}
$$

using L in (5.7). It is important to note that $\mathcal{L}$ is independent of π_i. In our problem, (5.16) can be further calculated as

$$\begin{cases} \mathcal{L}u_i - \pi_i(u_i)_y + \pi_i^2/4\kappa Q = 0 \\ \pi_{i+1}(z,t) = 2\kappa Q(D_y u_i), \end{cases} \tag{5.18}$$

where D_y denotes the partial differential operator with respect to y.

We already know from Section 3 that the solution of the semilinear PDE (3.2) exists uniquely with bounded spatial derivatives, so instead of $\mathcal{A}$ in (5.2), we can take a subset of $\mathcal{A}$ of which the controls are uniformly bounded. Also, note that $D_y u_i$ expresses the vega of the (approximated) market driver which we assumed to be a plain vanilla. As mentioned in Section 1, since plain vanilla options are positive vega products and $Q > 0$, we know that π_i is nonnegative from the definition in (5.18). If we assume that (2.2) is satisfied, then we see that the condition (4.5) is also satisfied. This precludes Y from becoming negative.

In order to apply the PIA, we need to check if the algorithm (5.18) satisfies the criteria of the PIA. The only criterion needed to be verified is the uniform Lipschitz condition on π_i. The following lemma proves this.

Lemma 5.2. $\{\pi_i\}_i$ *defined in* (5.18) *is uniformly Lipschitz continuous.*

Proof. From the Schauder estimate, we have

$$\|u_{i+1}\|_{2+\alpha} \leq C(\|g\|_{2+\alpha} + \|f^{\pi_n}\|_\alpha), \tag{5.19}$$

where C only depends on the Hölder norms of the coefficients of L^π, the domain Q_T, and ν_1 in (3.3). If g is continuous, we can approximate it uniformly in $2+\alpha$ norm by the Weierstrass approximation theorem as mentioned on page 71 in [6]. In our specific problem, $f^{\pi_n} = {\pi_n}^2/4\kappa Q$, so $\|f^{\pi_n}\|_\alpha$ is uniformly bounded thanks to the uniform boundedness of $\pi_i \in \mathcal{A}$. As the right hand side of (5.19) is uniformly bounded, $(u_i)_i$ is uniformly bounded in $2+\alpha$ norm, hence π_i is uniformly Lipschitz continuous from the second equation in (5.18). ♢

The PIA tells us that the u_i in (5.18) converges and the limit function is $V^{g,\mathcal{E}}$ which is $C^{2,1}$ and satisfies the HJB equation (4.4) in Q_T.

We will see later in the actual numerical example that the convergence to the solution happens fast. In the case of a plain call option as the market driver, we get a numerical solution very close to that of the semilinear PDE with only 1 iteration.

Proposition 5.3. $\|u_{i+2} - u_{i+1}\|_{2+\alpha} \leq C\kappa Q \|u_{i+1} - u_i\|_{2+\alpha}^2$

Proof. By definition and Proposition 5.1

$$\begin{cases} \mathcal{L}u_{i+2} - \pi_{i+2}(u_{i+2})_y + \pi_{i+2}^2/4\kappa Q = 0 \\ \mathcal{L}u_{i+1} - \pi_{i+1}(u_{i+1})_y + \pi_{i+1}^2/4\kappa Q = 0. \end{cases} \tag{5.20}$$

Subtracting these 2 equations and setting $v_{i+2} := u_{i+2} - u_{i+1}$,

$$\mathcal{L}v_{i+2} - \pi_{i+2}(v_{i+2})_y - (\pi_{i+2} - \pi_{i+1})^2/4\kappa Q = 0. \tag{5.21}$$

Since v_i is 0 on the parabolic boundary, from the Schauder estimate:

$$\|v_{i+2}\|_{2+\alpha} \leq C\|\frac{\pi_{i+2} - \pi_{i+1}}{4\kappa Q}\|_\alpha^2 = C\kappa Q\|(v_{i+1})_y\|_\alpha^2 \leq C\kappa Q\|v_{i+1}\|_{2+\alpha}^2 \tag{5.22}$$

♢

Proposition 5.3 shows that if the approximation of the solution is close enough to the classical solution of the semilinear PDE, $\{u_i\}_i$ converges quadratically to the solution. In other words, Proposition 5.3 shows the quadratic local convergence of the solutions of the PIA to the classical solution.

Corollary 5.4.

$$\|u_{i+1} - u_i\|_{2+\alpha} \leq (C\kappa Q\|u_1 - u_0\|_{2+\alpha})^{2^i-1}\|u_1 - u_0\|_{2+\alpha} \tag{5.23}$$

Proof. Use Proposition 5.3 and induction. ♢

6. Numerical Simulation

We now numerically investigate how the pricing and risks change with our model. We assume that a large amount of 2 year, 120 strike call is owned by investors outside the OTC market. We first price this structure using (2.8). Then, substituting this solution in (2.7), we price a different derivatives product, a 2 year, 100 strike call. We compare the results with the ones obtained from the Heston model. We used the explicit FDM method. We note that with sufficiently fine mesh in the discretization, the numerical solutions converge to the analytic ones. Hereafter, we refer to the 2 year, 120 strike call as 120 call and 2 year, 100 strike call as 100 call or at-the-money (ATM) call.

We use the parameters in Table 1.

Note that Feller's condition (2.2) is met and $Q > 0$. From Remark 2, the positive variance condition (2.6) is therefore satisfied.

We take our domain $\mathcal{E}$ to be a round rectangle and denote by S_{min}, S_{max}, v_{min}, and v_{max} the minimum and maximum values of the variables in the domain. In this example, we took $S_{min} = 0.5$, $S_{max} = 200$, $v_{min} = 0.00005$, and $v_{max} = 1.0$ and discretized each interval by 50. For the time interval $[0, 2]$, we discretized it similarly by 30,000. We denote by F_H the value F calculated in the Heston model and by F_N the value calculated in the new model. Similarly, we denote by V_H and V_N the corresponding values for an arbitrary V.

Parameter	Value
Q	0.0003
r	3.0%
ρ	-0.7571
η	0.3
ω	1.0
$\bar{v}$	0.04
κ	0.55

TABLE 1. Parameters for numerical simulation.

As in [7], for the calculation in the Heston model, we use the initial and boundary conditions:

$$
\begin{cases}
F_H(S,v,0) = \max(0,S-K)^+ & (S,v) \in \Omega \\
F_H(S_{min},v,t) = 0 & (S = S_{min}) \\
\frac{\partial F_H}{\partial S}(S_{max},v,t) = 1 & (S = S_{max}) \\
\frac{\partial F_H}{\partial t} - rS\frac{\partial F_H}{\partial S} + rF_H - \kappa\bar{v}\frac{\partial F_H}{\partial v} = 0 & (v = v_{min}) \\
F_H(S,v_{max},t) = S & (v = v_{max}).
\end{cases} \tag{6.1}
$$

The solution to the initial-boundary problem of Heston's PDE with conditions (6.1) is continuous up to the boundary, so we can use the value of F_H as the boundary condition for F_N. This way, the values of F_H and F_N match on the parabolic boundary.

We use corresponding boundary conditions for V.

With the parameters in Table 1, the drift in the second SDE of (2.4) is shifted by $\kappa Q\frac{\partial F_N}{\partial v}$, which in this case is calculated as $0.55 \times 0.0003 \times 77.188 = 0.0127$. This is about 58% of $\kappa\bar{v}$.

The result for the 120 call (which in our case is the market driver) is shown in Table 2.

Risks	Value	Delta	Vega	Vanna	Volga
Heston	2.6058	35.378%	70.940	3.8766	119.001
New Model	3.5121	42.457%	77.188	2.4132	-535.557

TABLE 2. Summary for 120 call at $S = 98.255$ and $v = 0.030049$

The result for the other derivative product (in our case, an at-the-money call) is shown in Table 3.

Risks	Value	Delta	Vega	Vanna	Volga
Heston	11.299	74.117%	79.238	-0.5666	-543.263
New Model	12.116	76.942%	78.824	-1.6201	-961.800

TABLE 3. Summary for at-the-money (ATM) call at $S = 98.255$ and $v = 0.030049$

The results are for $S = 98.255$ and $v = 0.030049$ at time $t = T = 2$. In volatility convention (i.e. standard deviation, as traders usually prefer this over variance), this value of v is equivalent to $\sigma = \sqrt{v} = 17.335\%$.

The obvious result is that the options are priced higher under the new model and we see it from Table 2 and Table 3. This is due to the current set-up that the 120 call (which is a positive vega product) is held outside of the OTC market. Since the OTC market is then overall short vega, or in other words, short volatility, the model correctly adjusts the level of the volatility which is now in demand. If we calculate the equivalent volatilities in the Heston model based on the prices we get from the new model, we get the correspondance shown in Table 4.

	120 Call	ATM Call
Heston	17.335%	17.335%
New Model	20.694%	20.090%
Difference	3.359%	2.755%

TABLE 4. Implied volatility calculated based on the risk calculated in the Heston model

From Table 4, we see that the volatility is higher, and the increments against the Heston volatilities are different for different structures. The result of Table 4 shows a skewness of the impact the market driver has on the volatility.

To understand how large this difference in the implied volatility is, we can assume that the vega traders maintain ranges between ±$10 million. With 3% difference in volatility as shown in Table 4, if they are short $10 million of vega, their mark-to-market loss would be -$30 million. If their goal is to raise $100 million of profit in a year, then this loss already corresponds to 30% of the annual target.

Figure 1 and Figure 2 show a simulation of the processes of the stock price and the volatility.

As mentioned in Section 1, this model prices-in not only the initial impact when some big position is traded with clients, but also the adjusted impact due to the change in the risk of the market driver. The risks change as the market moves, therefore the way traders hedge options changes under the new model. This is reflected in the graphs of the delta, vanna, and volga risks calculated in the new model compared to the ones calculated in the Heston model in Figure 3. The difference in each risk is plotted in Figure 4.

For example, when we check the delta on Table 2 and Table 3, the values are higher in the new model. This is because traders lose money when the stock price goes higher. To explain this in more detail, when the stock price goes higher, the vega of the 120 call gets larger since the stock price gets closer to the strike 120. This makes the traders in the OTC market get shorter in vega, hence they will even be more eager to buy the volatility in the market. This shifts the volatility higher. The consequence of this is that the traders will lose in mark-to-market because the value of the call they are short is greater now due to the spike in volatility. The new model anticipates this and asks the traders to buy more stocks beforehand so that they are hedged from this event.

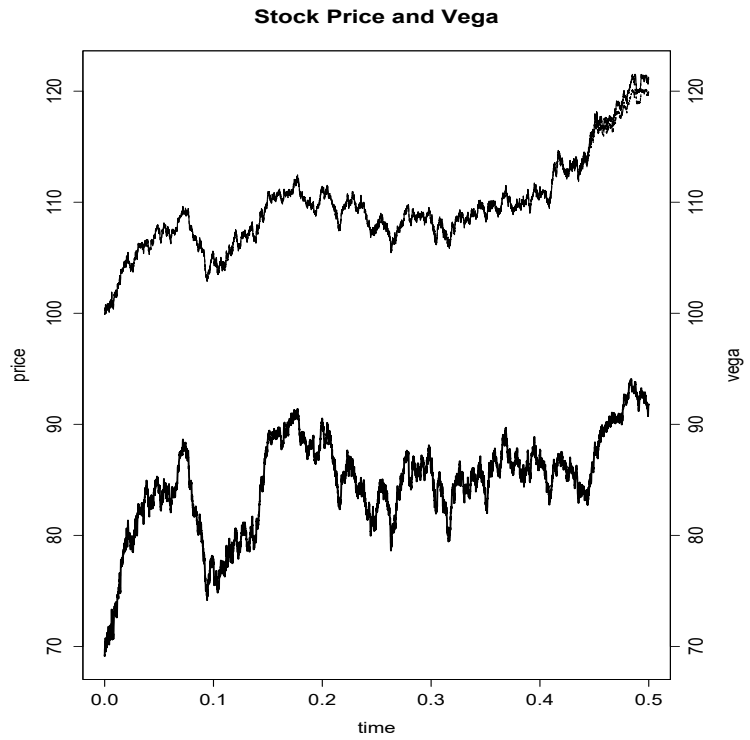

FIGURE 1. Simulation of the SDEs (2.4) for the first 6 months starting from $S = 100$ and $v = 0.04$ with F being the value of the 2Y 120 call. We plotted both the stock price processes of the Heston (dotted line) and of the new model (solid line) on the upper half of the graph. The largest difference in absolute value of the realizations of the two price processes is 0.7336, which corresponds to 73.36 basis points to the initial stock price. We used the drift $\mu = 0.05$. The lower graph shows how the vega of the call in the new model changes over time.

We now see what happens when we apply the PIA to the semilinear case in calculating the value of 120 call. We take $\pi_0 \equiv 0$ so that the solution of 0th iteration matches with the one from the Heston model. The result is shown in Figure 5. We tried up to 4th iteration as it implies convergence in numerical solution at this point as shown in Table 5.

Iteration	0th	1st	2nd	3rd	4th
Difference	2.3177	0.0322	7.72×10^{-6}	0.000	0.000

TABLE 5. Largest differences in absolute value between the numerical solutions of the approximated linear PDE and the original semilinear PDE. The figures could be regarded as the differences in percentage against the initial price of the stock as it is set to 100.

In Figure 6, we show a magnification of Figure 5 centered around the stock price where we saw the largest difference, which happened to be at-the-money.

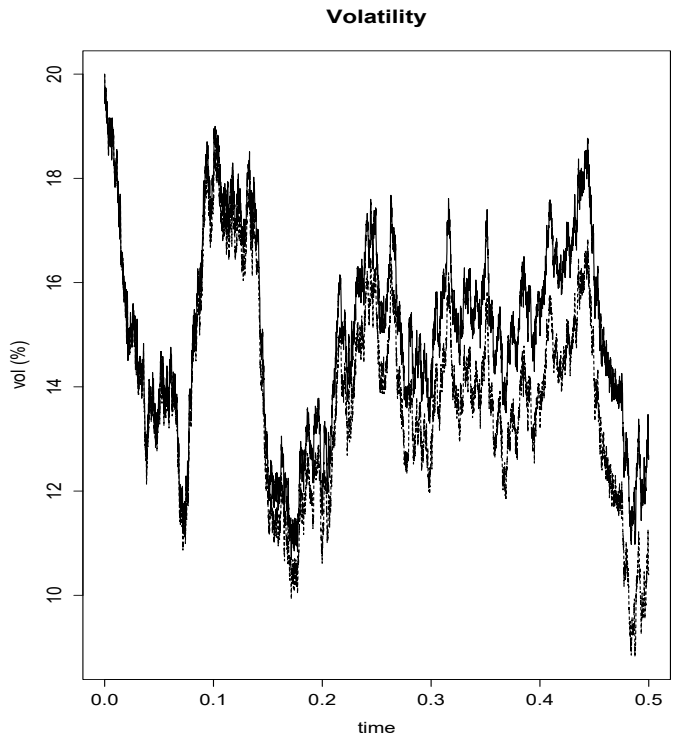

FIGURE 2. The volatility processes on the same simulation as in Figure 1, where the dotted line corresponds to that of the Heston model and the solid line to that of the new model.

We see in Figure 6 that the numerical solution of the semilinear PDE is different from that of the Heston model (0th iteration), but the 1st iteration in the PIA already brings the solution very close to that of the semilinear PDE. This is also implied by the result from Table 5. This means that the numerical solution of the semilinear PDE is well approximated by a series of linear PDEs. This is good news as we don't have to create a separate program to calculate the solution to the new model, but can just reuse the same program for the Heston model with modified coefficients. The PIA also appears to have better convergence compared to the explicit FDM on a Dirichlet boundary value problem of a second order semilinear elliptic PDE.

7. Conclusions

We introduced a new model which reflects the impact of a large position that is skewing the volatility market. We also introduced the Policy Improvement Algorithm. The algorithm lets us handle a semilinear PDE as a series of linear PDEs and at the same time keep the calculation load similar to that when we run the FDM on the original semilinear problem, thanks to the fast convergence of the iterations. This enables us to easily implement the new model in practice by reusing the resources used for the Heston model which has already been widely used in the industry.

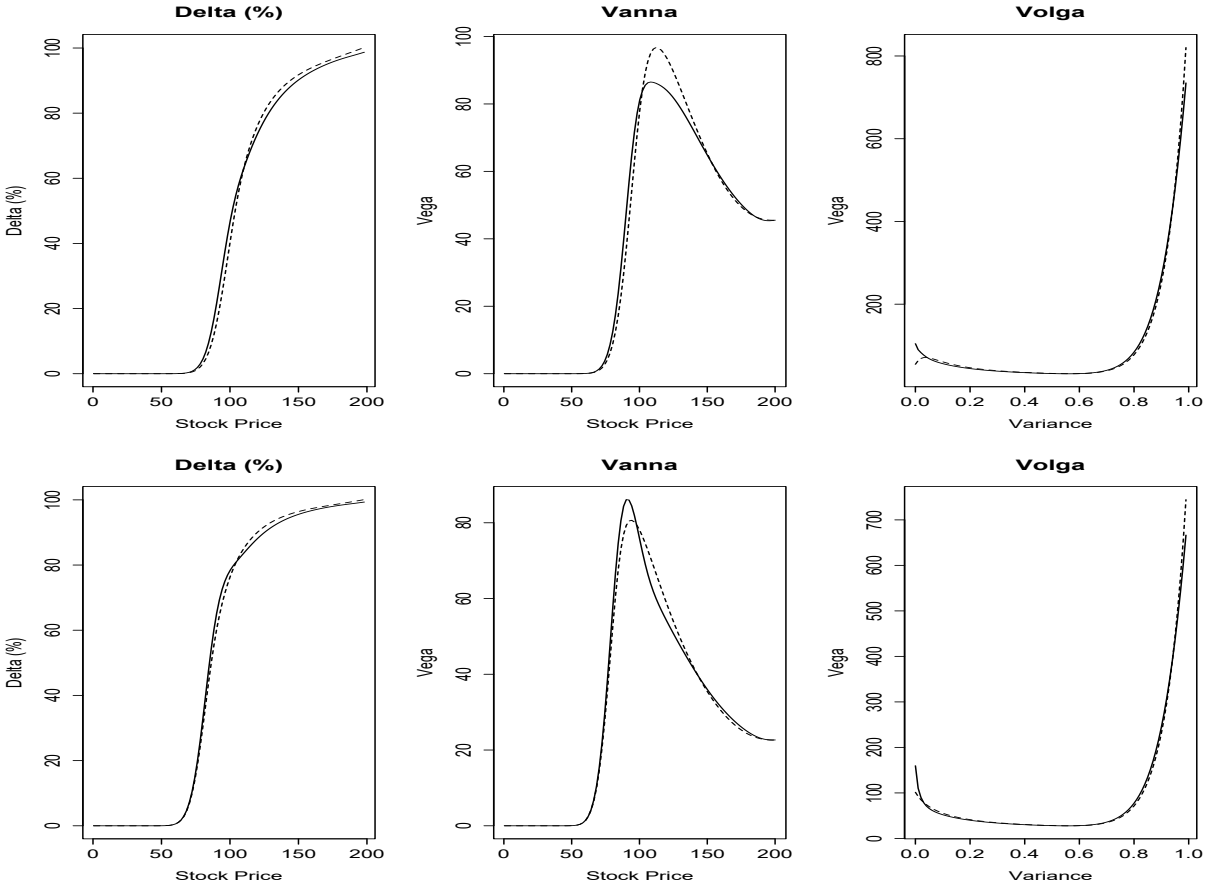

FIGURE 3. Risks of the calls; Top 3 charts are for 120 call and the bottom 3 are for ATM call. The solid lines indicate the risks calculated in the new model and the dotted line the corresponding risks calculated in the Heston model.

We only used a single product as a market driver, but we might try to extend this to the case when it is of a portfolio of several products. We only used a plain vanilla option as the market driver, but we should also be able to extend the model to be used for more exotic options. The difficulty then is to show the existence and uniqueness of the solution to the semilinear PDE (2.8) and to check if the solution satisfies the positive variance condition (2.6). If so, then by substituting this solution in the coefficient of the linear PDE (2.7), we can solve for the values of other derivatives products as in the case of the Heston model. It only takes relatively small effort to allow for the market asymmetry and to get the correct risks driven by the market driver.

The other difficulty in applying the model to actual trading appears in the calibration process. We assumed that we knew all the parameters including the detail of the market driver, but it may be challenging to recover these in the actual market, especially with more freedom in the model than in the Heston model and with limited market information.

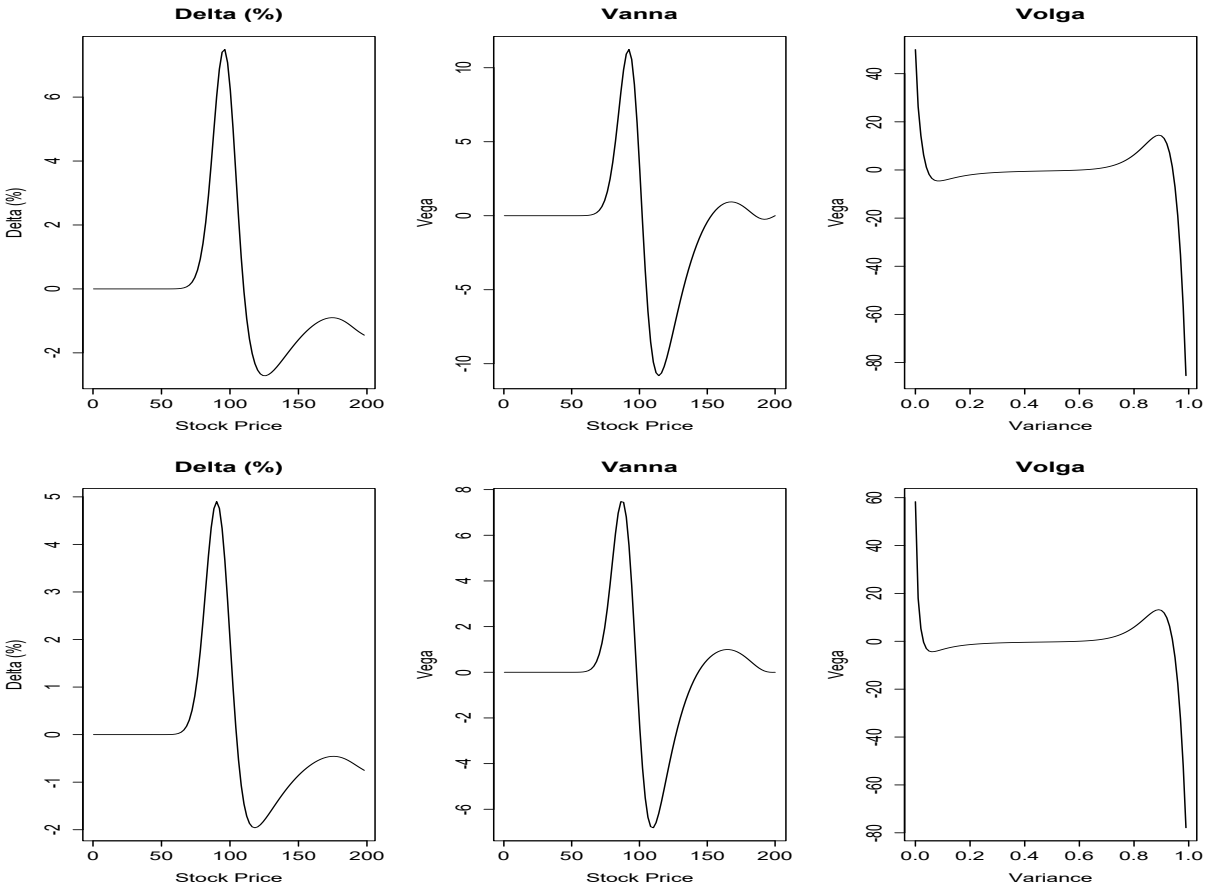

FIGURE 4. The difference plotted between the values in the new model and the Heston model from Figure 3.

ACKNOWLEDGEMENTS

We would like to thank Aleksandar Mijatović for reading the draft and providing helpful suggestions and insights.

REFERENCES

[1] BIS (Bank for International Settlements) 2015. Semiannual OTC Derivatives Statistics, 13 September 2015. http://www.bis.org/statistics/derstats.htm

[2] Duffy, D.J. 2006. *Finite Difference Methods in Financial Engineering: A Partial Differential Equation Approach*. John Wiley & Sons, Inc..

[3] Evans, L.C. 2010. *Partial Differential Equations*, American Mathematical Society, second edition.

[4] Fleming, W.H. and Soner, H.M. 1993. *Controlled Markov Processes and Viscosity Solutions*, Springer-Verlag.

[5] Frey, R. and Stremme, A. 1997. *Market Volatility and Feedback Effects from Dynamic Hedging*, Mathematical Finance, vol. 7(4), pp. 351-374.

[6] Friedman, A. 1964. *Partial Differential Equations of Parabolic Type*, Prentice-Hall.

[7] Heston, S.L. 1993. *A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options*, Review of Financial Studies, vol. 6(2), pp. 327-343.

[8] Hull, J.C. 2012. *Options, Futures, and Other Derivatives*. Pearson, eighth edition.

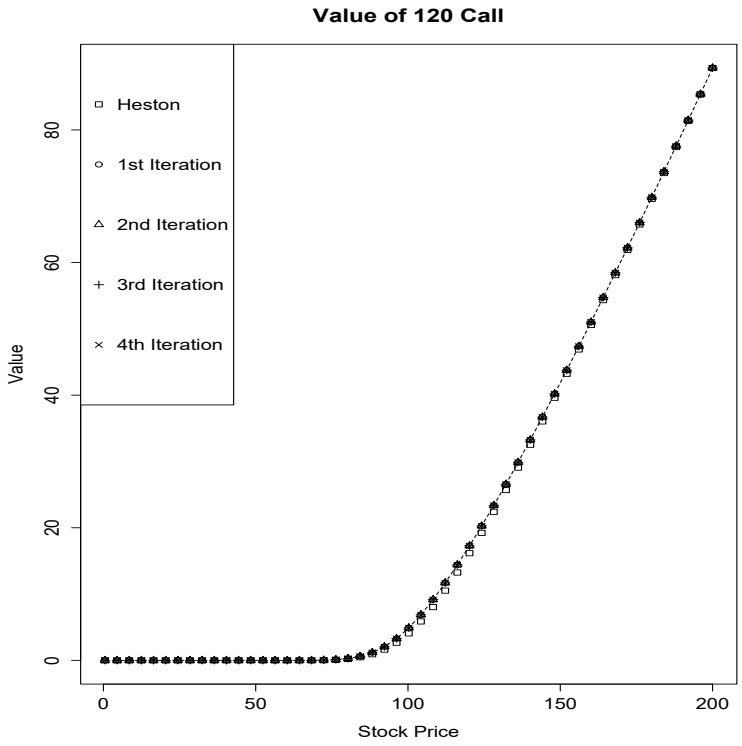

FIGURE 5. PIA results for 120 call. Dotted line is the solution using Finite Difference Method (FDM) directly on the semilinear PDE. v is taken as $v = 0.040048$.

[9] Ikeda, N. and Watanabe, S. 1981. *Stochastic Differential Equations and Diffusion Processes*, North Holland, North-Holland Mathematical Library.

[10] Jacka, S.D., Mijatović, A., and Širaj, D. *Policy Improvement Algorithm for Continuous Finite Horizon Problem*, forthcoming.

[11] Jacka, S.D., Mijatović, A., and Širaj, D. *Policy Improvement Algorithm for Controlled Multidimensional Diffusion Processes*, forthcoming.

[12] Karatzas, I. and Shreve, S.E. 1998. *Brownian Motion and Stochastic Calculus*, Springer Science+Business Media, Inc., second edition.

[13] Karlin, S. and Taylor, H.M. 1981. *A Second Course in Stochastic Processes*, Academic Press Inc.

[14] Ladyženskaja, O.A., Solonnikov, V.A., and Ural'ceva, N.N. 1968. *Linear and Quasi-linear Equations of Parabolic Type*, American Mathematical Society.

[15] Lieberman, G.M. 1996. *Second Order Parabolic Differential Equations*. World Scientific Publishing Co. Pte. Ltd.

[16] Lin, S. 2008. *Finite Difference Schemes for Heston Model*. Master's thesis.University of Oxford.

[17] Lindvall, T. and Rogers, L.C.G. 1986. *Coupling of Multidimensional Diffusions by Reflection*, The Annals of Probability, vol. 14(3), pp. 860-872.

[18] Lord, R., Koekkoek, R., and Dijk, D. 2010. *A Comparison of Biased Simulation Schemes for Stochastic Volatility Models*, Quantitative Finance, vol. 10(2), pp. 177-194.

[19] Platen, E. and Schweizer, M. 1998. *On Feedback Effects from Hedging Derivatives*, Mathematical Finance, vol. 8(1), pp. 67-84.

[20] Rockafellar, R.T. 1970. *Convex Analysis*, Princeton University Press.

[21] Sircar, K.R. and Papanicolaou, G. 1998. *General Black-Scholes Models Accounting for Increased Market Volatility from Hedging Strategies*, Applied Mathematical Finance, vol. 5(1), pp. 45-82.

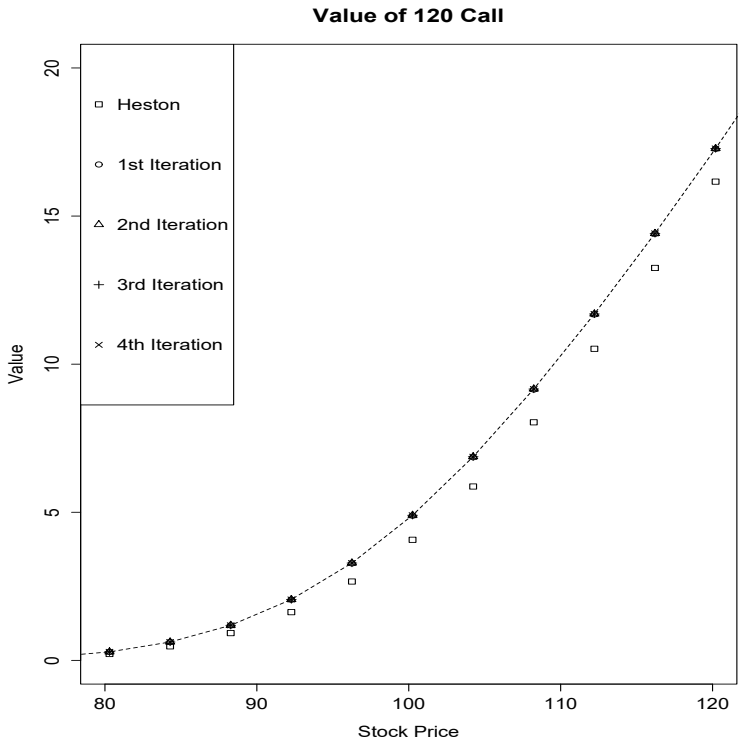

FIGURE 6. Magnification around at-the-money of Figure 5. We see that the first iteration already approximates well the numerical solution to the semilinear PDE.

[22] Smith, G.D. 1985. *Numerical Solution of Partial Differential Equations: Finite Difference Methods*. Oxford Applied Mathematics and Computing Science Series. Oxford, third edition.

[23] Soner, H. 2007. *Stochastic Representations for Nonlinear Parabolic PDEs* in *Handbook of Differential Equations*, Evolutionary Equations, vol. 3, Chapter 6, Elsevier B.V.

[24] Tavella, D., and Randall, C. 2000. *Pricing Financial Instruments: The Finite Difference Method*. John Wiley & Sons, Inc.

Appendix.

A. **Lemmas.** The lemmas stated here are more or less those in [10] and [11]. We only modify them to fit our problem. We show them here, however, so that this paper is self-contained.

A property that forms the basis of the following lemmas is that processes controlled by Markov policies are strong Markov processes (Theorem 4.20 in [12]).

Lemma A.1. *For every Markov policy π, $z \in \mathcal{E}$, $0 < t < T$, and any stopping time $\mathcal{S}$ that is almost surely less than $t \wedge \tau_\Omega$,*

$$\begin{aligned}\mathbb{E}\Big(\int_0^{t\wedge\tau} e^{-rs} f^\pi(Z_s^{z,\pi}, t-s)ds + e^{-r(t\wedge\tau)} g(Z_{t\wedge\tau}^{z,\pi}, t\wedge\tau)\Big| \mathcal{F}_{\mathcal{S}}\Big) \\ = \int_0^{\mathcal{S}} e^{-rs} f^\pi(Z_s^{z,\pi}, t-s)ds + e^{-r\mathcal{S}} V^{g,\mathcal{E},\pi}(Z_{\mathcal{S}}^{z,\pi}, t-\mathcal{S}).\end{aligned} \tag{A.1}$$

In particular, the process
$(\int_0^{T'} e^{-rs} f^\pi(Z_s^{z,\pi}, t-s)ds + e^{-rT'} V^{g,\mathcal{E},\pi}(Z_{T'}^{z,\pi}, T'))_{T'\le T}$ *is a uniformly integrable martingale.*

Proof. Let $\tau = \tau_{\mathcal{E}}(Z^{z,\pi})$ and $\tau_{\mathcal{S}} := \tau \circ \theta_{\mathcal{S}} = \tau_{\mathcal{E}}(Z^{z,\pi}_{\cdot+\mathcal{S}})$, where θ is the shift operator. Then $\tau_{\mathcal{S}} = \tau - \mathcal{S}$ holds almost surely, and we obtain

$$
\begin{aligned}
&\mathbb{E}\bigg(\int_0^{t\wedge\tau} e^{-rs} f^\pi(Z_s^{z,\pi}, t-s)ds + e^{-r(t\wedge\tau)} g(Z^{z,\pi}_{t\wedge\tau}, t\wedge\tau) \bigg| \mathcal{F}_{\mathcal{S}} \bigg) \\
&= \int_0^{\mathcal{S}} e^{-rs} f^\pi(Z_s^{z,\pi}, t-s)ds + \mathbb{E}\bigg(\int_{\mathcal{S}}^{t\wedge\tau} e^{-rs} f^\pi(Z_s^{z,\pi}, t-s)ds + e^{-r(t\wedge\tau)} g(Z^{z,\pi}_{t\wedge\tau}, t\wedge\tau) \bigg| \mathcal{F}_{\mathcal{S}} \bigg) \\
&= \int_0^{\mathcal{S}} e^{-rs} f^\pi(Z_s^{z,\pi}, t-s)ds + \mathbb{E}\bigg(\int_0^{t\wedge\tau-\mathcal{S}} e^{-r(s+\mathcal{S})} f^\pi(Z^{z,\pi}_{s+\mathcal{S}}, t-(s+\mathcal{S}))ds \\
&\qquad\qquad + e^{-r((t-\mathcal{S})\wedge(\tau-\mathcal{S})+\mathcal{S})} g\big(Z^{z,\pi}_{(t-\mathcal{S})\wedge(\tau-\mathcal{S})+\mathcal{S}}, (t-\mathcal{S})\wedge(\tau-\mathcal{S})+\mathcal{S},\big) \bigg| \mathcal{F}_{\mathcal{S}} \bigg) \\
&= \int_0^{\mathcal{S}} e^{-rs} f^\pi(Z_s^{z,\pi}, t-s)ds + e^{-r\mathcal{S}} \mathbb{E}\bigg(\int_0^{(t-\mathcal{S})\wedge\tau_{\mathcal{S}}} e^{-rs} f^\pi(Z^{z,\pi}_{s+\mathcal{S}}, t-(s+\mathcal{S}))ds \\
&\qquad\qquad + e^{-r((t-\mathcal{S})\wedge\tau_{\mathcal{S}})} g\big(Z^{z,\pi}_{(t-\mathcal{S})\wedge\tau_{\mathcal{S}}+\mathcal{S}}, (t-\mathcal{S})\wedge\tau_{\mathcal{S}}+\mathcal{S},\big) \bigg| \mathcal{F}_{\mathcal{S}} \bigg) \\
&= \int_0^{\mathcal{S}} e^{-rs} f^\pi(Z_s^{z,\pi}, t-s)ds + e^{-r\mathcal{S}} \mathbb{E}_x\bigg(\Big\{ \int_0^{(t-\mathcal{S})\wedge\tau} e^{-rs} f^\pi(Z_s^{z,\pi}, t-(s+\mathcal{S}))ds \\
&\qquad\qquad + e^{-r(t-\mathcal{S})\wedge\tau} g\big(Z^{z,\pi}_{(t-\mathcal{S})\wedge\tau}, (t-\mathcal{S})\wedge\tau\big) \Big\} \circ \theta_{\mathcal{S}} \bigg| \mathcal{F}_{\mathcal{S}} \bigg) \\
&= \int_0^{\mathcal{S}} e^{-rs} f^\pi(Z_s^{z,\pi}, t-s)ds + e^{-r\mathcal{S}} \mathbb{E}_{Z_{\mathcal{S}}^{z,\pi}}\bigg(\int_0^{(t-\mathcal{S})\wedge\tau} e^{-rs} f^\pi(Z_s^{z,\pi}, t-(s+\mathcal{S}))ds \\
&\qquad\qquad + e^{-r(t-\mathcal{S})\wedge\tau} g\big(Z^{z,\pi}_{(t-\mathcal{S})\wedge\tau}, (t-\mathcal{S})\wedge\tau\big) \bigg) \\
&= \int_0^{\mathcal{S}} e^{-rs} f^\pi(Z_s^{z,\pi}, t-s)ds + e^{-r\mathcal{S}} V^{g,\mathcal{E},\pi}(Z_{\mathcal{S}}^{z,\pi}, t-\mathcal{S})
\end{aligned}
$$

♢

By taking expectation on both sides of (A.1), we retrieve a corollary which is so-called Bellman's principle.

Corollary A.2. *For every Markov policy π, $z \in \mathcal{E}$, $0 < t < T$, and stopping time $\mathcal{S}$ which is almost surely less than or equal to $t \wedge \tau_\Omega$,*

$$
V^{g,\mathcal{E},\pi}(z,t) = \mathbb{E}\bigg(\int_0^{\mathcal{S}} e^{-rs} f^\pi(Z_s^{z,\pi}, t-s)ds \bigg) + e^{-r\mathcal{S}} \mathbb{E}(V^{g,\mathcal{E},\pi}(Z_{\mathcal{S}}^{z,\pi}, t-\mathcal{S})). \tag{A.2}
$$

We now use the method of mirror coupling [17].

Lemma A.3. *For every Lipschitz Markov control and small enough $\epsilon > 0$, there exists $\delta > 0$ such that the following holds for every $z_1, z_2 \in \mathcal{E}$: if $\|z_1 - z_2\| < \delta$ then there exist processes $\tilde{Z}^{z_1,\pi}$ and $\tilde{Z}^{z_2,\pi}$ that have the same laws as $Z^{z_1,\pi}$ and $Z^{z_2,\pi}$ respectively such that*

$$\|\tilde{Z}_t^{z_1,\pi} - \tilde{Z}_t^{z_2,\pi}\| \le G_{\tau_t} \quad on \quad t < \rho_\delta$$

and

$$\tilde{Z}_t^{z_1,\pi} = \tilde{Z}_t^{z_2,\pi} \quad on \quad t \ge \rho_0$$

for every $t \ge 0$, where

$$\rho_c := \inf\left\{t \ge 0; \|\tilde{Z}^{z_1,\pi} - \tilde{Z}^{z_2,\pi}\| = c\right\} \quad , \quad (\inf \phi = \infty)$$

for any $c \ge 0$, G is the squared Bessel process of dimension $1+\epsilon$ started at $\|z_1 - z_2\|$, and $(\tau_t)_{t\ge 0}$ is a stochastic time change with the property

$$\tau_t \le \frac{t}{\nu_1}, \quad t \ge 0.$$

For the proof of Lemma A.3, we refer to [11].

Lemma A.4. *For every Lipschitz Markov policy π, the function $V^{g,\mathcal{E},\pi}(\cdot, t)$ is continuous with bounded initial condition.*

Proof. Let $\epsilon > 0$ and $\hat{\delta} \le \delta$. For $\|z_1 - z_2\| \le \hat{\delta}$, we calculate $|V^{g,\mathcal{E},\pi}(z_1,t) - V^{g,\mathcal{E},\pi}(z_2,t)|$.

$$\begin{aligned}
&|V^{g,\mathcal{E},\pi}(z_1,t) - V^{g,\mathcal{E},\pi}(z_2,t)| \\
&= \left|\mathbb{E}\left(\int_0^{t\wedge\tau_{z_1}} e^{-rs} f^\pi(\tilde{Z}_s^{z_1,\pi}, t-s)ds + e^{-r(t\wedge\tau_{z_1})} g(\tilde{Z}_{t\wedge\tau_{z_1}}^{z_1,\pi}, t\wedge\tau_{z_1})\right)\right. \\
&\qquad \left. - \mathbb{E}\left(\int_0^{t\wedge\tau_{z_2}} e^{-rs} f^\pi(\tilde{Z}_s^{z_2,\pi}, t-s)ds + e^{-r(t\wedge\tau_{z_2})} g(\tilde{Z}_{t\wedge\tau_{z_2}}^{z_2,\pi}, t\wedge\tau_{z_2})\right)\right| \\
&< \left|\mathbb{E}\left(\int_0^{\rho_0} e^{-rs}\left\{f^\pi(\tilde{Z}_s^{z_1,\pi}, t-s) - f^\pi(\tilde{Z}_s^{z_2,\pi}, t-s)\right\}ds \middle| I_{\rho_0\le\rho_\delta}\right)\right| \\
&+ \left|\mathbb{E}\left(\int_{\rho_0}^{t} e^{-rs}\left\{f^\pi(\tilde{Z}_s^{z_1,\pi}, t-s) - f^\pi(\tilde{Z}_s^{z_2,\pi}, t-s)\right\}ds\right.\right. \\
&\qquad \left.\left. + e^{-rt}\left\{g(\tilde{Z}_t^{z_1,\pi}, t) - g(\tilde{Z}_t^{z_2,\pi}, t)\right\} \middle| I_{\rho_0\le\rho_\delta}\right)\right| \\
&+ \left|\mathbb{E}\left(\int_0^{t} e^{-rs}\left\{f^\pi(\tilde{Z}_s^{z_1,\pi}, t-s) - f^\pi(\tilde{Z}_s^{z_2,\pi}, t-s)\right\}ds\right.\right. \\
&\qquad \left.\left. + e^{-rt}\left\{g(\tilde{Z}_t^{z_1,\pi}, t) - g(\tilde{Z}_t^{z_2,\pi}, t)\right\} \middle| I_{\rho_0>\rho_\delta}\right)\right| \\
&= B_1 + B_2 + B_3.
\end{aligned} \tag{A.3}$$

For B_1, since f^π is Lipschitz continuous, we can take $\delta_1 \in (0,\delta)$ small enough such that

$$B_1 < C\|\tilde{Z}_s^{z_1,\pi} - \tilde{Z}_s^{z_2,\pi}\| < \epsilon/2. \tag{A.4}$$

For B_2, due to the definition of $\tilde{Z}$, the processes $\tilde{Z}_t^{z_1,\pi}$ and $\tilde{Z}_t^{z_2,\pi}$ take the same values in this time frame in consideration, so $B_2 = 0$.

Due to the boundedness of f^π and g, the last term B_3 could be bounded by some constant multiplied by $\mathbb{P}(\rho_0 > \rho_\delta)$. If we denote by $\rho_\delta(\mathcal{Y})$ and $\rho_0(\mathcal{Y})$ the first hitting times of the levels δ' and 0 respectively for any process $\mathcal{Y}$, we have from Lemma A.3

$$\mathbb{P}(\rho_\delta < \rho_0) \leq \mathbb{P}\Big(\rho_\delta(G_\tau) < \rho_0(G_\tau)\Big) \leq \mathbb{P}\bigg(\rho_\delta\Big(G_{\frac{1}{\nu_1}}\Big) < \rho_0\Big(G_{\frac{1}{\nu_1}}\Big)\bigg). \tag{A.5}$$

Using the scale property of the squared Bessel process we get

$$\mathbb{P}\bigg(\rho_\delta\Big(G_{\frac{1}{\nu_1}}\Big) < \rho_0\Big(G_{\frac{1}{\nu_1}}\Big)\bigg) = \mathbb{P}\bigg(\rho_\delta\Big(\frac{1}{\nu_1}G\Big) < \rho_0\Big(\frac{1}{\nu_1}G\Big)\bigg) = \mathbb{P}(\rho_{\nu_1\delta}(G) < \rho_0(G)). \tag{A.6}$$

Recall that the scale function of the Bessel process with dimension $1+\epsilon$ is given by $s(z) := z^{\frac{1-\epsilon}{2}}$, and that the process G starts at $\|z_1 - z_2\| < \hat{\delta}$. Hence we obtain

$$\mathbb{P}(\rho_{\nu_1\delta}(G) < \rho_0(G)) = \frac{s(\|z_1 - z_2\|) - s(0)}{s(\nu_1\delta)} \leq \left(\frac{\hat{\delta}}{\nu_1\delta}\right)^{\frac{1-\epsilon}{2}}. \tag{A.7}$$

We set $\delta = \delta_1$ and take $\hat{\delta} \in (0, \delta)$ small enough so that

$$2C\left(\frac{\hat{\delta}}{\nu_1\delta}\right)^{\frac{1-\epsilon}{2}} < \frac{\epsilon}{2}. \tag{A.8}$$

Collecting what we calculated, we have proved that $\|z_1 - z_2\| < \hat{\delta}$ implies $|V^{g,\mathcal{E},\pi}(z_1, t) - V^{g,\mathcal{E},\pi}(z_2, t)| < \epsilon$, so we have uniform continuity of $V^{g,\mathcal{E},\pi}(\cdot, t)$. ♢

Lemma A.5. *For every Lipschitz Markov policy π, the function $V^{g,\mathcal{E},\pi}$ is continuous.*

Proof. If we proved the continuity of $V^{g,\mathcal{E},\pi}$ with respect to t for fixed z, the statement is proved using the triangle inequality and Lemma A.4. Therefore, we prove the continuity in t with fixed z. Due to Corollary A.2, we have

$$\begin{aligned} V^{g,\mathcal{E},\pi}(z, t+\delta) - V^{g,\mathcal{E},\pi}(z, t) = \mathbb{E}\bigg(\int_0^\delta e^{-rs} f^\pi(Z_s^{z,\pi}, t-s)ds\bigg) \\ + e^{-r\delta}\mathbb{E}\Big(V^{g,\mathcal{E},\pi}(Z_\delta^{z,\pi}, t) - e^{r\delta}V^{g,\mathcal{E},\pi}(z, t)\Big). \end{aligned} \tag{A.9}$$

Applying Lemma A.4, we obtain

$$|V^{g,\mathcal{E},\pi}(z, t+\delta) - V^{g,\mathcal{E},\pi}(z, t)| \leq C\delta + C'\mathbb{E}(\|Z_\delta^{z,\pi} - z\|). \tag{A.10}$$

The SDE for $Z_\delta^{z,\pi}$ yields

$$Z_\delta^{z,\pi} - z = \int_0^\delta \mu_\pi(Z_s^{z,\pi}, s)ds + \int_0^\delta \sigma_\pi(Z_s^{z,\pi}, s)dW_s. \tag{A.11}$$

Therefore, we have

$$\begin{aligned} |V^{g,\mathcal{E},\pi}(z, t+\delta) - V^{g,\mathcal{E},\pi}(z,t)| &\le C\delta \\ &+ C'\mathbb{E}\bigg(\bigg|\int_0^\delta \mu_\pi(Z_s^{z,\pi}, s)ds\bigg|\bigg) + C''\mathbb{E}\bigg(\bigg|\int_0^\delta \sigma_\pi(Z_s^{z,\pi}, s)dW_s\bigg|\bigg). \end{aligned} \tag{A.12}$$

The second term on RHS can be bounded by some multiple of δ as μ_π is bounded. For the last term, using Jensen's inequality and Burkholder-Davis-Gundy inequality,

$$\begin{aligned} \mathbb{E}\bigg(\bigg|\int_0^\delta \sigma_\pi(Z_s^{z,\pi}, s)dW_s\bigg|\bigg) &\le \bigg(\mathbb{E}\bigg(\int_0^\delta \sigma_\pi(Z_s^{z,\pi}, s)dW_s\bigg)^2\bigg)^{\frac{1}{2}} \\ &\lesssim \bigg(\mathbb{E}\bigg(\int_0^\delta \sigma_\pi^2(Z_\delta^{z,\pi}, s)ds\bigg)\bigg)^{\frac{1}{2}}. \end{aligned} \tag{A.13}$$

This proves the continuity of $V^{g,\mathcal{E},\pi}$ with respect to t with fixed z. Therefore, the continuity of $V^{g,\mathcal{E},\pi}$ is proved. ♢

University of Warwick, Coventry, United Kingdom

E-mail address: J.Maeda@warwick.ac.uk

University of Warwick, Coventry, United Kingdom

E-mail address: S.D.Jacka@warwick.ac.uk